What Would a Cuttlefish Do?

A prediction waiting to be tested — March 2026

Waldemar Lis — The Self-Experiencing Universe

archive.org/details/@not_foryou

The Experiment That Has Not Been Done

In 2013, researchers demonstrated that honeybees can step back from problems they cannot solve. Given a visual discrimination task and an option to opt out, bees chose to withdraw from difficult or impossible choices far more often than from easy ones. They appeared, in some meaningful sense, to know when they did not know.

No one has run the equivalent experiment on a cuttlefish.

This is a gap in the literature, but it is not a permanent one. The experiment is feasible, the species is available, and the question is interesting enough that someone will eventually ask it. The question this simulation asks is: before that experiment is run, what does the framework predict?

Reasoning from What We Know

The previous simulation in this series established that individual honeybees occupy a specific position in VCP space — characterized by moderate self-modeling capacity, as measured by the opt-out experiment. Sepia officinalis was placed higher on the self-modeling axis, based on behavioral observations compiled by Professor Jennifer Mather: the cuttlefish monitors arms it cannot see, substitutes one arm for another before attempting an action that would require it, and performs a characteristic head bob that appears to function as an uncertainty check before committing to a camouflage decision.

If self-modeling capacity is what the opt-out experiment measures — and that is the framework's core assumption — then Sepia should perform differently than a bee on such a test. Not just somewhat differently. Substantially differently.

The Prediction

The framework predicts that on difficult trials — trials designed to be at the edge of discriminability — Sepia officinalis would opt out approximately 84% of the

time, compared to the bee's 52%. On easy trials, the prediction is a modest increase in baseline caution. On impossible trials, near-complete withdrawal.

More importantly: the framework predicts that Sepia would generalize the opt-out strategy to entirely novel tasks at a substantially higher rate than bees — approximately 57% of individuals, compared to 35% for bees. This generalization is what distinguishes a true self-model from learned avoidance of specific difficult stimuli. A system that genuinely models its own uncertainty applies that model to new situations. A system that merely learns to avoid hard patterns does not.

Sensitivity — the sharpness of the discrimination between easy and impossible — is predicted to be near its theoretical maximum for Sepia. The cuttlefish is not expected to be uncertain about whether it is uncertain.

What This Would Mean

If the experiment is run and the results fall within the predicted range, the triangulation would be complete: bee, swarm, and cuttlefish would occupy three distinct regions of the consciousness geometry, each separated in a way that is not merely quantitative but structural.

If the results fall outside the predicted range — if Sepia performs no better than a bee, or performs worse — the framework would need revision. The self-modeling axis would need to be reconsidered, or the behavioral markers used to position Sepia on it would need to be reinterpreted.

Either outcome is informative. That is what makes it a prediction rather than a description.

Status

This simulation is the second in the VCP empirical series. It follows the Perry & Barron mapping (Simulation 1, March 2026) and precedes any response from Professor Lars Chittka (Queen Mary University of London), whose interpretation of the bee metacognition data may require recalibration of the anchor point. If that recalibration occurs, a post-Chittka version of this prediction will be issued.

Full technical documentation is available to researchers through archive.org/details/@not_foryou.